

EXHIBIT 5

HOW GOOGLE ANONYMIZES DATA

Anonymization is a data processing technique that removes or modifies personally identifiable information; it results in anonymized data that cannot be associated with any one individual. It's also a critical component of Google's commitment to privacy.

By analyzing anonymized data, we are able to build safe and valuable products and features, like autocompletion of an entered search query, and better detect security threats, like phishing and malware sites, all while protecting user identities. We can also safely share anonymized data externally, making it useful for others without putting the privacy of our users at risk.

Two of the techniques we use to protect your data

Generalizing the data

There are certain data elements that are more easily connected to certain individuals. In order to protect those individuals, we use generalization to remove a portion of the data or replace some part of it with a common value. For example, we may use generalization to replace segments of all area codes or phone numbers with the same sequence of numbers.

Generalization allows us to achieve k anonymity, an industry standard term used to describe a technique for hiding the identity of individuals in a group of similar persons. In k anonymity, the k is a number that represents the size of a group. If for any individual in the data set, there are at least k-1 individuals who have the same properties, then we have achieved k-anonymity for the data set. For example, imagine a certain data set where k equals 50 and the property is zip code. If we look at any person within that data set, we will always find 49 others with the same zip code. Therefore, we would not be able to identify any one person from just their zip code.

If all individuals in a data set share the same value of a sensitive attribute, sensitive information may be revealed simply by knowing these individuals are part of the data set in question. To mitigate this risk, we may leverage l-diversity, an industry-standard term used to describe some level of diversity in the sensitive values. For example, imagine a group of people searched for the same sensitive health topic (e.g. flu symptoms) all at the same time. If we look at this data set, we wouldn't be able to tell who searched for the topic, thanks to k-anonymity. However, there may still be a privacy concern since everyone shares a sensitive attribute (i.e. the topic of the query). L-diversity means the anonymized data set would not only contain flu searches. Rather, it could include other searches alongside the flu searches to further protect user privacy.

Adding noise to data

Differential privacy (also an industry-standard term) describes a technique for adding mathematical noise to data. With differential privacy, it's difficult to ascertain whether any one individual is part of a data set because the output of a given algorithm will essentially appear the same, regardless of whether any one individual's information is included or omitted. For example, imagine we are measuring the overall trend in searches for flu across a geographic region. To achieve differential privacy, we add noise to the data set. This means we may add or subtract the number of people searching for flu in a given neighborhood, but doing so would not affect our measurement of the trend across the broader geographic region. It's also important to note that adding noise to a data set may render it less useful.

Anonymization is just one process we use to maintain our commitment to user privacy. Other processes include strict controls on user data access, policies to control and limit joining of data sets that may identify users, and the centralized review of anonymization and data governance strategies to ensure a consistent level of protection across all of Google.